

Allocation of radio resources in a CDMA2000 cellular system

FIELD OF THE INVENTION

5 This invention relates to data transmission in a CDMA cellular telephone system and in particular – to allocation of transmission resources of a transmitter among connected subscribers.

BACKGROUND OF THE INVENTION

Cellular telephone systems are increasingly used for transmission of data from
10 various sources to subscribers. The data are usually requested by a subscriber and are usually transmitted to the cellular system from appropriate sources over the Internet, then transmitted through the core network and a radio access network (RAN) of the cellular system to the base station, with which the subscriber unit is connected, and finally radio-transmitted from the base station to the subscriber unit.
15 In conformance with the Internet Protocol (IP), all data are transmitted asynchronously as packets. Data arriving at the base station are accumulated in buffer storage, where a storage bin is designated to each actively connected subscriber unit. The present invention is concerned with the final lag of transmission, namely from a buffer storage bin, over the base station's radio
20 transmitter, to the subscriber's receiver.

Cellular data transmission differs from cellular voice transmission in several aspects: (a) The required rate of transmission, in the case of data, varies greatly with time, as well as among subscriber units, while for voice it remains constant. (b) Moreover, the maximum required rate, for data, may exceed that for voice by
25 orders of magnitude. (c) Unlike voice, data need not, generally, be transmitted continuously, but may be transmitted in bursts (i.e. many packets in close succession), with considerable intervals between them; however certain types of data applications (notably streaming types) have different tolerances to lengths of intervals. (d) Various subscribers may be given different degrees of quality of

service (QoS) – for example, in terms of guaranteed minimum transmission rate. In order to accommodate these characteristics, new operating standards are being introduced to cellular systems. In particular, systems using the Code Division Multiple Access (CDMA) mode of transmission, have a new standard, known as
5 CDMA 2000.

As is known, a channel in a CDMA system is defined by a particular code out of a set of N mutually orthogonal codes, known as a Walsh set. According to the CDMA 2000 standard there are defined, for any one radio transmission facility (e.g. a radio carrier), a set of $N=2^n$ fundamental channels, in terms of N Walsh codes,
10 where N is typically 64 ($n=6$). Typically, a fundamental channel (FCH, also known as 1x channel, which is similar to a regular voice channel) carries data at a rate of about 10 Kbits per second. By defining suitable common subsets of the codes, fundamental channels are combinable, in a hierarchical manner into higher-rate (i.e. wider-bandwidth) channels as follows: A set of $N/2$ (e.g. 32) 2x channels, each
15 carrying about 20 Kbits per second; a set of $N/4$ (e.g. 16) 4x channels, each carrying about 40 Kbits per second; and so on. The levels are also designated by a parameter k , whose value is 0 for the lowest level (i.e. for FCH), 1 for the second level, 2 for the third level and so on. The transmission rate (TR) at any level is then simply that of the lowest level (e.g. 10 Kb/s) multiplied by 2^k . The corresponding
20 Walsh codes are designed as a hierarchical binary structure, illustrated in Figure 1, wherein, at each level, a code is a subset of certain two codes at the next lower level. Thus, for example, a 4x code is a subset of four related fundamental channel codes (two levels lower). The relationship between channels at the various levels, as determined by the code structure, is sometimes expressed in genealogical terms,
25 such as “father”, “brother”, “son”, “grandson”, etc. When any particular wide channel is used, all related channels at the lower levels (i.e. all the “offsprings”), as well as all the directly related channels at the higher levels (i.e. all the “ancestors” in direct lineage) are not available; it is noted, though, that other related channels (such as “brothers” and “uncles” and their offsprings) are available, if not otherwise
30 allocated. At any time, any available channel at any level may be allocable to any

active subscriber unit, subject to the hierarchy discussed above and to certain constraints discussed below.

In common with other cellular systems, in a system operating under the CDMA 2000 standard, the signal power transmitted to any subscriber unit (SU) is a function of the transmission conditions (which depend, *inter alia*, on the distance between the subscriber unit and the base station), whereby the power is adjusted to maintain a given received signal-to-noise ratio. The signal-to-noise ratio is, however, also a function of the channel bandwidth and thus of the rate level of the channel; the higher the rate, and thus the wider the band, the higher the noise level and thus the higher the required transmission power. It is also noted that the total power of all signals transmitted at any instant is subject to a maximum value, characteristic of the transmitter.

Periodically channels must be allocated to subscriber units in order to carry to them data, addressed to them that have accumulated in the buffer storage. The duration of the period is in the order of a few hundred milliseconds. According to current practice, allocation is carried out, by a so-called scheduling engine, near the beginning of each period typically as follows: A storage bin is selected at random and the amount of data accumulated therein is compared with a series of threshold values; according to the outcome, a commensurate transmission rate, if any, is determined. If a channel of that rate is available, it is allocated to the subscriber unit corresponding to the bin, provided that the power required to transmit it to the subscriber unit does not cause the total power to exceed the maximum. Failing this, a channel of half the determined rate, if available, is allocated to that subscriber unit, again subject to the power test; and so on. Next, another bin is randomly chosen and the same process is followed. This cycle is repeated until there are no channels left or until any allocation would cause the maximum available power to be exceeded or until there is no data waiting in storage. Variations of this procedure are also in common practice.

Allocation procedures in present practice, as outlined above, do not optimally utilize the limited transmission resources, which are the overall data rate capacity and the maximum overall power. In particular, they allow spending an inappropriately large portion of the power on high rate transmission to subscriber
5 units having poor radio reception; they also cause transmission rates to be dependent solely on buffered data sizes, which by themselves are random, and possibly on random selection. Moreover, procedures in present practice do not generally include QoS considerations and also cannot be geared to any business policy, such as would control transmission so as to maximize some variable (which
10 may, for example, be overall transmission rate or overall revenue).

SUMMARY OF THE INVENTION

The present invention is of a method, and an apparatus to carry out the method, to allocate transmission resources of a cellular transmitter among active subscriber units (SUs) in communication therewith, for which there is data stored in an
15 associated buffer storage, the allocation being so as to attempt to maximize one or more given variables according to any defined business policy. As discussed above in the Background section, the given variable may be overall transmission rate (throughput) or overall revenue (according to some billing model) or adherence to one or more promised QoS attributes, such as average rate QoS. The basic idea of
20 the method is to calculate, prior to the beginning of each period, one or more priority factors for each active subscriber unit and to allocate channels according to the priority factors, commencing with the one having highest value. The calculation formula operates on the given variable and on the level of power required in each case; in general, the priority factor is inversely related to the required power. The
25 allocation procedure also considers all the relevant operational constraints, such as minimum and maximum amounts of stored data per subscriber unit, the hierarchical structure of the channels and the maximum overall power, as well as constraints dictated by the standard, by fairness considerations and by service

agreements. The latter may include components such as a minimum average rate to each subscriber unit or a maximum delay between data packages.

In the Detailed Description below are explained exemplary algorithms to implement the method of the invention, for three specific business policies; the
5 choice of the particular business policies is also itself exemplary, intended to illustrate the principles of the invention. These algorithms are summarized below, but neither they nor the choice of the particular business policies addressed by them should be construed as limiting the scope of the invention.

A business policy implemented in the first exemplary algorithm is one that
10 aims simply at maximizing throughput, i.e. making the total amount of data transmitted over any period of time as large as possible. To implement this policy, priority factors are calculated for each active subscriber unit and for every possible level of transmission rate (i.e. channel bandwidth). The rate level is constrained by the availability of channels and, for each subscriber unit, by the current queue size
15 (i.e. the amount of data waiting in the respective buffer bin). Each priority factor (PF) is calculated as the ratio of some function of the respective bit rate to the power required for transmission to the subscriber unit over a fundamental channel. Allocation for any period is carried out in a recursive procedure, whereby first all possible priority factors are calculated and the one with the highest value selected;
20 the corresponding subscriber unit is assigned a channel at the corresponding rate level, if consistent with constraints imposed by available system resources; the system resources are then reevaluated and the cycle is repeated with the next highest priority factor.

A business policy implemented in the second exemplary algorithm is one that
25 aims at maximizing throughput while assuring a minimal average data rate during any one call. This policy is implemented as follows: For each active subscriber unit, a discrepancy factor is calculated, as the ratio of the promised rate to the average rate of actual transmission since the beginning of the call, and then a priority factor is calculated, as the ratio of a function of the discrepancy factor to the power

required for transmission to the subscriber unit over a fundamental channel. Allocation for any period is, again, carried out in a recursive procedure, whereby first all possible priority factors are calculated and the one with the highest value selected; the corresponding subscriber unit is assigned a channel at the highest
5 possible rate level, if consistent with constraints imposed by available system resources; the system resources are then reevaluated and the cycle is repeated.

A business policy implemented in the third exemplary algorithm is one that aims at maximizing revenue. Such a policy must be based on a selectable billing model. Typically a billing model is a formula that relates the charge for a call to the
10 amount of data transmitted and to the QoS (e.g. average rate of transmission). As in the implementation of the first policy, discussed above, the periodic allocation procedure is recursive, whereby at each cycle, a priority factor is calculated for each as yet unallocated subscriber unit and for every possible rate level (under the same constraints) and the highest priority factor indicates the next subscriber unit
15 and the rate level of the channel to be allocable to it. However, in this case each priority factor is calculated as follows: First an incremental revenue of the respective subscriber unit is calculated as the revenue to be generated by transmission to the subscriber unit at the respective rate over the coming time period; the calculation is based on the billing model. Next, a net revenue gain
20 (NRG) is calculated, as the difference between the respective incremental revenue and the loss in revenue that would be incurred by all other subscriber units that are not allocated a channel. Finally, the priority factor is calculated as the ratio of the respective net revenue gain to the respective required power.

BRIEF DESCRIPTION OF THE DRAWINGS

25 In order to understand the invention and to see how it may be carried out in practice, a preferred embodiment will now be described, by way of non-limiting example only, with reference to the accompanying drawings, in which:

Fig. 1 is a schematic diagram depicting the hierarchical structure of channels at various rate levels in a cellular data transmission system with CDMA 2000 standard.

Fig. 2 depicts a table of priority factors, used in some of the embodiments of
5 the invention.

Fig. 3 is a flow diagram showing the principal operations required for channel allocation according to a first embodiment of the invention.

Fig. 4 is a flow diagram showing the principal operations required for channel allocation according to a second embodiment of the invention.

10 **Fig. 5** is a flow diagram showing the principal operations required for channel allocation according to a third embodiment of the invention.

Fig. 6 is a schematic block diagram illustrating apparatus to carry out the methods of the invention.

DETAILED DESCRIPTION OF THE INVENTION

15 Figure 2 shows a two-dimensional table, in which rows correspond to active subscriber units and columns – to the various channel rate designations (e.g. 1x, 2x, 4x, etc.). The entries in the table are priority factor values, denoted in the figure by “PF” with the row- and column indices; thus $PF_{j,k}$ is the value of the priority factor for subscriber unit j when receiving data over a channel of rate level k . The priority
20 factors are calculated and used in ways specific to the various embodiments, explained below. It is noted that, in general, a priority factor is calculated as the ratio of some specific variable to a power-related variable. The latter is generally some monotonically increasing function of the power required to transmit data to the respective subscriber unit at some given signal-to-noise ratio and some given
25 rate level; usually the rate level for this calculation is the lowest, i.e. that of a fundamental channel, and the function is preferably linear (or simply an identity). The reason for the power-related variable being in the denominator of the priority

factor formula is that allocating a channel to a subscriber unit with relatively high power requirement is wasteful of the limited power resource of the transmitter.

A **first** embodiment of the method of the invention is directed to allocating channels according to a business policy that seeks to **maximize throughput**, i.e. the average overall rate of data flow to subscriber units. Equation (1) shows the formula for calculating the priority factor under this policy.

$$(1) \quad PF_{j,k} = \frac{(TR_{FCH}).2^k}{P_j^m}$$

where:

- 10 $PF_{j,k}$ = Priority factor of subscriber j at rate level k
 TR_{FCH} = Transmission rate of a fundamental channel (FCH) (say 10kb/s)
 P_j = Power required for transmitting fundamental channel to subscriber j
 m = constant.

- 15 The specific variable in the numerator of the formula is simply equal to the transmission rate (TR) over the respective channel, i.e. that of the corresponding rate level (appearing as column heads in the table of Fig. 2). For example, for a fundamental (1x or FCH) level ($k=0$) the rate is 10 Kbits/sec., for a 2x level ($k=1$) – 20 Kbits/sec., for a 4x level ($k=2$) – 40 Kbits/sec., etc.; the numerator may also be
20 any variable proportional to the transmission rate. As explained above for the general case, the denominator consists of P raised to some power m , where P is equal or proportional to the power required to transmit data to the respective subscriber unit over, say, a 1x channel (FCH).

- Fig. 3 is a flow diagram showing the principal operations carried out by the
25 allocation procedure according to a first embodiment of the invention. The procedure, which is iterative, is executed prior to the beginning of each period (of,

say, 200 milliseconds) and its outcome is the allocation of channels to subscriber units for transmission over that period. During the procedure, a list of available channels, initially obtained, for example, from a Walsh codes manager (see Fig. 6), is maintained and updated upon each allocation iteration. The list is subject to the hierarchical structure of the rate levels, as in Fig. 1; this means that allocation of a particular channel at a certain rate level eliminates from the list not only that channel, but also all corresponding channels at the lower levels (i.e. those directly or indirectly branching from it), as well as corresponding channels at all the higher levels (i.e. those from which it directly or indirectly branches). For example, with reference to Fig. 1, if the channel represented by the rightmost circle at level 4x (k=2) were allocated, then also the channels represented by the rightmost two circles at level 2x (k=1), as well by the rightmost four circles at level 1x (k=0), would become unavailable; additionally, the channels represented by the rightmost circle at each of levels 8x (k=3), 16x (k=4) and 32x (k=5), as well as the channel at the top level (k=6), would become unavailable.

First (21), priority factors are compiled, in a table such as in Fig. 2, for active subscriber units (in rows) and for relevant rate levels (in columns). Of the subscriber units, there are listed only those for which there are data waiting in the buffer storage. For each listed subscriber unit, priority factors are calculated, according to Equation (1), for every rate level up to a maximum level. The maximum level is determined as that which does not exceed the lower of (a) the amount of respective data in the buffer, divided by the duration of a period; and (b) the maximum rate currently available. It should be noted that the currently available maximum rate is identical for all subscriber units and is determined by scanning the available channels register; initially, it is that of the highest level in the system (e.g. 16x). All entries whose required power exceeds the available power are eliminated from further consideration (22).

All priority factors in the table are then compared with each other (23) and the one with the highest value is selected. The respective subscriber unit is then

allocated a channel at the respective rate level and is eliminated from the table(24). If there still remain listed other subscriber units, the table is updated (22) as follows: (a) The available channels register is updated, a new maximum available rate level is determined accordingly and entries in all columns of the table above
5 that rate level are eliminated; (b) the remaining available power is calculated, as that which was previously available (which is maximum power during the first iteration) minus that required for the transmission just allocated, and all entries for which the required power would exceed the available power are eliminated.

The remaining entries are again compared (23) and the process is cyclically
10 iterated until there are no entries left in the table. This will normally occur when there are no available channels left or when the maximum available power level has been reached or when there are no data left in the buffer storage. At this point the allocation information is fed to the transmitter and transmission of data to the selected subscriber units over the allocated channels may proceed over the ensuing
15 period. Toward the end of that period a similar allocation procedure is initiated and the whole process repeats as required.

A second configuration of the method of the invention is directed to allocating channels according to a business policy that seeks to maximize throughput while maintaining promised minimum average rates. A minimum average rate is a value
20 that may be promised to a subscriber, usually as part of a service level agreement (SLA). In this configuration the table of priority factors has a single column and a priority factor is calculated for each active subscriber unit, using Equation (2):

$$(2) \quad PF_j = \frac{(TR_{promised} / TR_{average})_j}{P_j^m}$$

where:

25 PF_j = Priority factor of subscriber j

$TR_{promised}$ = Transmission rate promised to subscriber j

$$\begin{aligned} TR_{average} &= \text{Average transmission rate to subscriber } j \text{ since beginning of} \\ &\quad \text{message} \\ m &= \text{constant} \end{aligned}$$

It will be seen that Equation (2) is similar to Equation (1) except for the
5 numerator, calculated as follows: The amount of data transmitted to each subscriber
unit, as well as the time elapsed, from the beginning of the respective current data
transfer process, are continuously accumulated and their ratio is calculated, to
become the average data rate. The specific variable in the numerator of the priority
factor formula is calculated as the ratio of the promised minimum rate to the last
10 calculated average data rate. The specific variable is thus a measure of the
discrepancy between the actual average rate received by a subscriber unit and the
rate promised thereto. As before, this variable is divided by P, the required
transmission power at the lowest rate level (1x or FCH), with exponent (i.e. raised
to some power) m.

15 Fig. 4 is a flow diagram showing the principal operations carried out by the
allocation procedure according to a second embodiment of the invention. After
calculating a table of priority factors (41), as explained above, the procedure is,
again, iterative and, in general, similar to the procedure of the first configuration,
except as follows: The subscriber unit with the highest priority factor value is
20 selected (42) and an attempt is made to allocate a channel to it; the chosen channel
is of the highest rate level available, provided that the corresponding required
power does not exceed the available transmitter power (after accounting for
previously allocated channels and subscriber units – which are null at the first
iteration). If (43) the power test fails, and if (44) the attempted rate is still wider
25 than 1x, an attempt is made (45) to allocate a channel at a just lower rate level. The
power test is then repeated (43) and the loop iterated until either the power test is
satisfied, in which case the subscriber unit is allocated a channel at the last
attempted rate level and (4-6) the list of available channels and the available power
are updated, as in the first configuration, or (44) the attempted allocation is of a

fundamental channel (1x) and fails the power test. In either case, the respective subscriber unit is eliminated from further consideration. As before, the procedure is iterated as long as (4-8) there are unallocated subscriber units in the table or the transmitter resources (power or channels) are exhausted.

5 A **third configuration** of the method of the invention is directed to allocating channels according to a business policy that seeks to **maximize total revenue**. It is similar to the first configuration, differing from it mainly in the calculation of the priority factors. Two formulas, shown as Equations (3) and (4) serve for this calculation. The first formula, shown in Equation (3), for calculating revenue,
10 expresses a revenue function that reflects a billing model, that is – the manner of charging a subscriber for data transmission.

$$(3) \quad R_{j,t} = B_{j,t} [C_0 + a \cdot \log(1 + b \cdot (TR_{average})_{j,t})]$$

where:

15 $R_{j,t}$ = Revenue accumulated from subscriber j at a time t (since call initiation)

$B_{j,t}$ = Total number of bits received by subscriber j at a time t (since call initiation)

$(TR_{average})_{j,t}$ = Average transmission rate to subscriber j until time t (since call initiation)

20 C_0, a, b = constants

 The method of this invention, in general, and the third embodiment, in particular, contemplates any billing model. The formula in Equation (3) is exemplary and expresses the revenue function of a possible billing model; any other model and corresponding function and formula could be equally used. The
25 revenue function of Equation (3) includes a component that is related to the amount of data transmitted and a component that is related to the average rate of transmission. The second component represents a premium for fast rates, which is

one of the attributes of quality of service (QoS). Other attributes of QoS that may be considered in a billing model are average delay, average jitter and reliability (the latter being a function of signal-to-noise ratio and hence of transmitted signal power). Inclusion of particular QoS attributes in the billing model may be determined as part of the service level agreement (SLA), differentially for each subscriber unit. Moreover, the attributes may be applied for any subscriber unit differentially as a function of the indicated application for (or type of) the data. It will be appreciated that while the invention is described with reference to a specific revenue function and corresponding billing calculation, as well as specific data collection for billing, these are exemplary and the invention contemplates use of any billing model and a corresponding revenue function in the channel allocation procedure. The exemplary revenue formula shown in Equation (3) stipulates that the revenue be proportional to the amount of data transmitted and to a logarithmic function of the average rate of transmission. The latter variable is determined in a manner similar to that in the second embodiment, as explained above.

Equation (4) shows a formula for calculating priority factors as function of anticipated revenue.

$$(4) \quad PF_{j,k} = \frac{(NRG)_{j,k}}{P_{j,k}^m}$$

where:

$(NRG)_{j,k}$ = Net Revenue Gain when transmitting to user j a channel of rate k

$P_{j,k}$ = Power required for transmitting a channel of rate k to a user j

m = constant

The specific variable in the numerator of Equation (4) is the net revenue gain (NRG) for the corresponding subscriber unit and the corresponding rate level. This variable, NRG, is calculated, on the basis of any billing model, such as that expressed by equation (3) for example, as follows: First, a “revenue gain” is

calculated, as the revenue that would be generated by sending data to the subscriber unit at the indicated channel rate level over the coming period; this is done by using the revenue formula twice – first for the revenue generated by the subscriber unit from the beginning of the current data transfer up to the present, then – for that generated up to the end of the coming period (whereby the value of the average transmitted rate is updated) and by taking the difference between the two results. Next, a “loss of alternative revenue” is calculated by applying a differential calculation, similar to the above, to each of the other subscriber units, but assuming that no data will be transmitted during the coming period. It is to be noted that the result will usually be negative, due to the reduction in the average rate value. Finally the “loss of alternative revenue” values of all other subscriber units (which, as noted, are usually negative) are added to the “revenue gain” of the candidate entry, to obtain the net revenue gain, NRG. The latter is then divided by a function (e.g. power m) of the respective required transmitting power (similarly to the previously described embodiments), to obtain the corresponding priority factor.

The allocation procedure in this embodiment may be according to any one of two preferred algorithms. The first algorithm is generally similar to that of the first configuration, as illustrated by the flow chart of Fig. 3 and as described above. The second algorithm is illustrated by the flow chart of Fig. 5; it is, again, basically similar to the first algorithm, except as follows: After selecting (51) from the priority factor table the entry with the highest value (which entry, to be termed “candidate entry”, is associated with a certain “candidate subscriber unit” and a certain “candidate rate level”), the table is checked (52) whether there exists for the candidate subscriber unit an entry at a level just one higher than the candidate rate level, and whether a father channel is available ; if both these conditions are met, the table is further checked (53) as to whether a brother channel to the candidate channel is available and whether there exists at least one other entry at the candidate level. If not all of these conditions are met, the candidate subscriber unit is allocated a channel at the candidate level (54) and the cycle is repeated as in the

other algorithms; otherwise, allocation is made as follows: The priority factors of all the other entries at the candidate level (if more than one) are compared (55) and the highest one (associated with a “candidate brother subscriber unit”) is added (56) to the priority factor of the candidate entry, to obtain a combined priority factor.

5 The combined priority factor is then compared (57) with the priority factor of the higher-level entry; if the latter is larger, the candidate subscriber unit is allocated (58) the higher-level channel; otherwise, both the candidate subscriber unit and the candidate brother subscriber unit are allocated (59) respective brother channels at the candidate level. The cycle is then repeated as in the other algorithms, until there

10 are no relevant entries left in the table (60).

Fig. 6 depicts schematically a preferred apparatus for carrying out the method according to the invention. The apparatus comprises a scheduler 10, that preferably resides on the premises of a corresponding Radio Access Network 30 (see below). Scheduler 10 includes an Allocator 16, a buffer storage 12 and a flow controller 14.

15 Allocator 16 consists of a suitably programmed digital processor, which is designed to operate in real time, that is – to carry out the computations required for allocating all available channels within a time duration that is considerably shorter than a period (in the cycle of allocation and transmission). Traffic data intended for active subscriber units is received from a core network 20 in the cellular system and

20 stored in buffer storage 12 – in FIFO bins corresponding to the subscriber units. Information about stored quantities in the storage bins is fed to Allocator 16. The allocator also receives, from a Walsh Codes Manager 32, information about the availability of the Walsh codes (i.e. the available channels at the various rate levels) and, from a Radio Channel Quality Indicator 34, information about the radio

25 reception quality (hence – the required power) for the various active subscriber units. Both Walsh Codes Manager 32 and Radio Channel Quality Indicator 34 are part of the cellular system. Allocator 16 has stored within it a computer program to carry out the calculations and procedures, according to any embodiment of the method of the invention, as described above, as well as all relevant constants and

parameters, including those related to business policies. The latter may also be fed to the Allocator from some management module within the cellular system.

Allocator 16 periodically carries out channel allocation, according to any of the above-described methods of the invention, and feeds the results to Flow Control
5 Module 14. The Flow Control Module 14, which may be a suitably programmed digital processor, accordingly retrieves appropriate amounts of data from relevant bins in the buffer storage 12 and sends them, together with the corresponding channel identities, to Radio Access Network 30 (which is part of the cellular system), to be radio transmitted.

10 The apparatus of Fig. 6 may be realized as an add-on module, to be connected to a cellular system, or as an inherent module that is integrated with the cellular system. Embodiments other than that of Fig. 6 are also possible within the scope of the invention. In particular, the scheduler, as a whole, or any of its components, such as the allocator, the buffer storage and the flow controller, can be realized
15 within, or as part of, equipment within the cellular system that carries out channel allocation and flow control also by other methods or serves for other functions within the cellular system. It is only essential that one or more digital processors be programmed to carry out the functions of channel allocation and data flow control according to any of the methods of the invention.

20 Although the apparatus and the method of the invention have been described in terms of channels allocation for one transmitter (with respect to subscriber units in communication with it), it will be understood that the method is obviously also applicable, and any one apparatus may be operative, to simultaneously allocate channels for a plurality of transmitters. In particular, the apparatus, preferably
25 being logically and physically associated with a Radio Access Network (RAN), is preferably operative to allocate channels for each of the transmitters served by the RAN.

Furthermore, although the method and apparatus of the invention have been described above in terms of a cellular telephone transmission system, it will be

understood that they are also applicable to other transmission systems having data transmission characteristics such as addressed above. It will also be understood that any processor in apparatus according to the invention may be a suitably programmed computer. Likewise, the invention contemplates a computer program
5 being readable by a computer for executing the method of the invention. The invention further contemplates a machine-readable memory tangibly embodying a program of instructions executable by the machine for executing the method of the invention. It will further be understood that other configurations and embodiments of the method and apparatus of the invention are possible, all coming within the
10 scope of the invention.

In the method claims that follow, alphabetic characters and Roman numerals used to designate different operations are provided for convenience only and do not imply any particular order of performing the operations.